

MULTI-DISTRICT PREFERENCE MODELLING

GEOFFREY PRITCHARD AND MARK C. WILSON

ABSTRACT. Generating realistic artificial preference distributions is an important part of any simulation analysis of electoral systems. While this has been discussed in some detail in the context of a single electoral district, many electoral systems of interest are based on multiple districts. Neither treating preferences between districts as independent nor ignoring the district structure yields satisfactory results. We present a model based on an extension of the classic Eggenberger-Pólya urn, in which each district is represented by an urn and there is correlation between urns. We show in detail that this procedure has a small number of tunable parameters, is computationally efficient, and produces “realistic-looking” distributions. We intend to use it in further studies of electoral systems.

1. INTRODUCTION

In order to test the average-case performance of electoral systems, or to test models of electoral phenomena such as inter-election swing, it is useful to be able to generate artificial preference distributions for a society.

We are particularly interested in multi-winner elections of parliamentary type, in which the electorate is partitioned into districts, one or more representatives are chosen from each district, and these representatives together make up the parliament. It is crucial to consider the variability between districts, as otherwise we obtain very unrealistic results. For example, if simulating plurality voting in single-winner districts (“First Past the Post”) where Party A has nationally 55% of the vote, if districts are identical then Party A will win 100% of seats. At another extreme, if all voters within a district behave identically then Party A will win 55% of the seats (ignoring integer rounding issues). Also, the vote share of Party A will be exactly 55% in each district under the first assumption, while under the second, in 55% of districts it achieves 100% of the vote. However, real-world FPTP elections are not consistent with these assumptions. The way districts are drawn typically assures each of the major parties dominance in some number of seats, with a smaller number of marginal ones, and so party A will typically have a vote share distribution across districts that looks very different from either of these. For a real-world example, see Figure 1.

In social choice theory (basically corresponding to the single-district case), the theoretical performance of voting systems, for example with respect to manipulability, has been studied in detail using various preference distributions [BL94; Reg+06]. This kind of work may be useful for studying, for example, party list proportional systems. However, we are not aware of any work for district-based electoral systems. In this article we present a general method for generating artificial voting data in districts which has a small number of tunable parameters, is computationally efficient, and produces “realistic-looking” distributions. We cover only the case of plurality ballots.

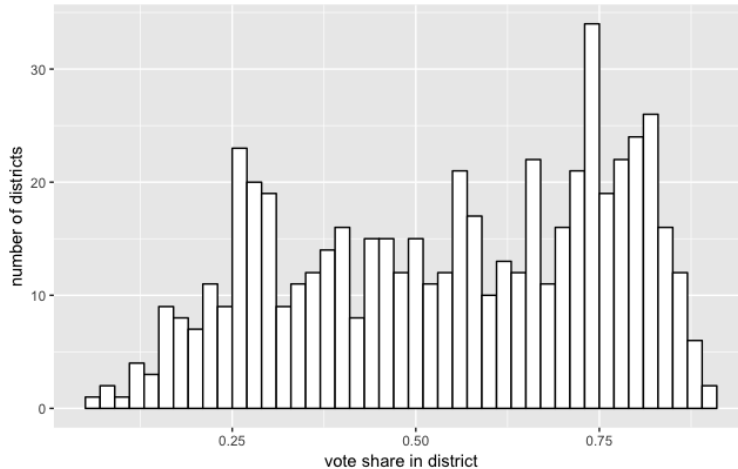
The rest of the paper is as follows. In Section 2 we introduce the basic probabilistic models that we will use for generating artificial data, and discuss their theoretical properties.

Key words and phrases. simulation, Pólya urn model, electoral systems.

Supported by the NZ Marsden Fund.

The authors acknowledge the support of the New Zealand Marsden Fund. The authors thank B. Grofman for useful discussions.

FIGURE 1. UK election 2015, Labour party 2-party percentage in England & Wales



In Section 3 we report on detailed experiments and explain the effect of varying parameters in the model. Most of this involves the case of two parties, but we also go beyond this in Section 3.3. Section 4 shows how the model can cast light on phenomena such as the cube law for the vote-seat mapping and theories of inter-election swing. In Section 5 we discuss possible future work.

2. URN MODELS

2.1. Eggenberger-Pólya urn models. We review the basic abstract framework, which in its essentials goes back to Eggenberger and Pólya in [EP23], although they in turn drew on the work of still earlier authors. Let K be a nonnegative integer and \mathcal{C} a finite set of *colors*. Consider a container (*urn*) into which we place initially some number a_i of indistinguishable *balls* of each color $i \in \mathcal{C}$, such that the total number of balls is finite. At each step we randomly choose a ball from the urn, and replace it together with K balls of the same color. When $K = 0$, this is simply sampling with replacement, but for larger values of K , the number of balls increases at each iteration with a bias toward colors that already occur frequently. The initial distribution of balls of course influences the distribution at later times. Most obviously, no color not already present in the initial distribution can ever be present in the urn.

This process can be used to model contagion and imitation. It has been widely used in the electoral context (see, e.g. [BL94]), in which the balls represent voters, and the colors candidates (or parties, or choices at a plebiscite, etc.). The urn model can then be interpreted as describing what happens when social influence occurs. In particular when $K = 1$, we can interpret it as adding new voters one at a time, and each voter imitating the vote of a randomly selected previous voter. The case $K = 1$ can be used to generate data from the so-called Impartial Anonymous Culture distribution of preferences [BL94]. Adding voters can be interpreted as meaning that eligible but previously uncommitted voters make a choice, or that new voters become eligible, for example by attaining the minimum voting age.

2.2. Extension to multiple urns. We extend the basic model above to one with N urns, to be thought of as representing electoral districts. Let p be a probability, that is, a real number between 0 and 1 (inclusive). In each urn there is an initial distribution of balls as above — any possible overall distribution between urns and within urns is possible. At each step, we first choose an urn u uniformly at random. With probability p , we choose to

draw a ball from some other urn (chosen uniformly at random from among all urns other than u); otherwise we draw a ball from u . When the ball has been drawn, it is returned to the urn from whence it came, and an additional ball of the same color is added to u .

2.3. Theoretical properties of the model. The known theoretical properties of the single Eggenberger-Pólya urn can provide some insight into the behaviour of the more complex multiple-urn model. Suppose a single urn initially contains a_1, \dots, a_m balls of colors $1, \dots, m$ respectively, and each ball drawn is replaced along with one more ball of the same color. Let X_1, X_2, \dots be random variables giving the colors of the balls drawn, and Y_1, Y_2, \dots be random m -vectors giving the proportions of the colors after each draw (*i.e.* $Y_n(i)$ is the proportion of the balls that have color i after the first n draws). Then

- (1) With probability 1, the sequence (Y_n) converges to a limit Y ;
- (2) Y is a random variable with the Dirichlet(a_1, \dots, a_m) distribution;
- (3) Conditional on Y , the random variables X_1, X_2, \dots are independent with probability distribution Y .

(See [BM73].) If there are only two colors, we may consider Y_n to be scalar-valued (since then $Y_n(2)$ is simply $1 - Y_n(1)$, adding no further information); in this case the limit Y has a Beta(a_1, a_2) distribution.

A particularly simple special case occurs when $a_1 = \dots = a_m = 1$ *i.e.* the urn begins with one ball of each color. The limiting Dirichlet($1, \dots, 1$) distribution is then the uniform distribution on the m -simplex; for the two-color case this is the uniform distribution on the interval $[0, 1]$.

Turning to multiple urns but remaining in the case where the initial numbers of balls in each urn are equal, we can make some observations about what to expect. Firstly, since each step adds a ball to one randomly selected urn, the urns will likely end up containing different numbers of balls. However, once the number n of steps taken has become large, there will be similar numbers of balls in each urn (the differences between the numbers of balls in different urns will be on the order of \sqrt{n} , but the numbers themselves are of order n).

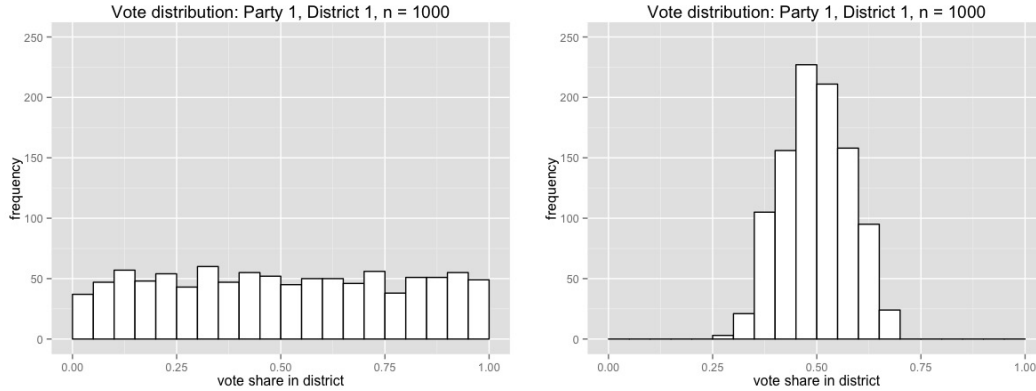
Secondly, when $p = 0$, the urns do not interact. Each urn behaves as a single Eggenberger-Pólya urn, independently of the others.

Thirdly, when $p = 1$ and N is not small, the behaviour will be close to that exhibited in the case where there is a single large urn, leading to substantial correlation between the ball counts in urns, for each color.

3. EVALUATION

From now on we use the motivating interpretation of the above models, namely plurality voting in districts. In this case each district corresponds to a unique urn, each ball represents a voter, and each color represents a candidate. The probability p is a parameter measuring the likelihood of inter-district imitation by voters, while K is a parameter controlling the strength of imitation, whether within a district or between two districts. The urn models are dynamic, and in principle could be used to model the evolution of voting behavior over time. We adopt a static interpretation here — we think of an urn model with given parameters as a stochastic process that we run for a long time until the situation is almost stable and we have a total number of voters suitable for the situation under study, and then use the resulting vote distribution as a simplified description of an artificial society.

The results above show that we expect the districts to have roughly equal numbers of voters at all times during the process, assuming that we start with equal numbers of voters

FIGURE 2. District vote share distribution when $p = 0$ (L) and $p = 0.2$ (R)

in each district. We now want to explore the effect on the final state of the process of varying p, K and the initial distribution of voter preferences across parties in each district.

In order to test the realism of the model, we carried out extensive simulation (C++ program available at [give URL according to publisher instructions](#)). The baseline computations involved $N = 100$, and we also performed some for $N = 1000$. We concentrated on values of m from 2 to 4, and varied p from 0 to 1 in steps of 0.1. We experimented with various initial conditions. For each case (value of N, m, p and initial conditions) we ran 1000 simulations. We mostly used $K = 1$ but also considered some higher values of K .

We display the results using several types of graphs. For party A we fix a district and show a histogram of the vote share in that district. Note that this is similar to what is shown in Figure 1, but different in that it shows vote shares in a single district across many replications of the experiment, rather than vote shares across all districts for a single replication.

We then show histograms of the party's total vote share V and the number of seats S that it wins under FPP. We also compute the correlation between A 's share of the total vote in the first 50 (arbitrarily labelled "north") districts and in the second 50 (labelled "south") districts. We then give a scatterplot of the mapping from V to S .

3.1. Two parties. We fix $m = 2$.

3.1.1. Symmetric initial conditions. Here we started with one voter of each party in each district. When $p = 0$, the distribution of the vote share of Party 1 in District 1 is shown in the left picture in Figure 2 (as predicted by the theory, the urn process produces something drawn from a uniform distribution) and the distribution of this party's seat share in the left picture of Figure 4 (as predicted by the theory, the urn process produces something drawn from a binomial distribution). The total vote share of Party 1 across all districts is shown in the left panel of Figure 3. As expected, we see no correlation between the party's vote in the northern and southern districts — see the left picture in Figure 5. The left picture in Figure 6 shows a positive association between votes and seats, again as expected.

Increasing the value of p to 0.2 leads to quite different results (see the right-hand pictures in Figures 2-6). The distribution of the party's vote in a given district becomes more concentrated; the influence of other districts now makes it much less likely that this district will deliver a result wildly different from the national average. At the same time, the seat distribution becomes *less* concentrated. The stabilization of the within-district vote puts many more districts in play, while inter-district imitation has a self-reinforcing effect: a high vote in one district not only wins that seat, but improves the party's prospects in







FIGURE 3. Total vote share distribution when $p = 0$ (L) and $p = 0.2$ (R)FIGURE 4. Seat share distribution when $p = 0$ (L) and $p = 0.2$ (R)FIGURE 5. North-south correlation when $p = 0$ (L) and $p = 0.2$ (R)FIGURE 6. Vote - seat mapping when $p = 0$ (L) and $p = 0.2$ (R)

FIGURE 7. The case $p = 1$: seats (L) and seats vs votes (R)FIGURE 8. District vote (L), total vote (M) and seat (R) distribution when $p = 0$: initial conditions 2 : 2.

every other district as well. Another view of this effect is provided by the north-south correlation graph (Figure 5). Also, the mapping of votes to seats has a larger slope, a point we shall return to in Section 4.1. Note that the distribution of total vote share changes very little, because changing the value of p has large district level effects (whether a new voter imitates a voter in the same district) but total vote share depends only on whether they imitate *some* voter.

As p increases, the trend observed above continues. Figure 7 shows the behaviour in the extreme case $p = 1$: the entire electorate is now behaving (approximately) like a single well-mixed district, and it is quite likely that all of the parliamentary seats will be won by the same party. (Due to the symmetry of the initial conditions, each of our two parties has the same probability of achieving such a total victory.) Note that the vote shares (both nationally and in each district) are still quite close to 50-50; it is the lack of diversity among districts that makes such lopsided seat-shares possible. This case could be thought of as resembling a single district with an initial content of 50 voters of each type.

We repeated the above experiments for the case where each party initially has 2 voters in each district (Figures 8 and 9). When $p = 0$, the district vote distribution is Beta(2,2) (density proportional to $x(1-x)$); consequently, the total vote is a little less variable than in Figure 3. The seat-share distribution is still binomial, unchanged from Figure 4, and the north-south correlation is still zero (graph omitted). The seat-vote relationship is shown in Figure 10. Introducing inter-district mimicry again has the effect of homogenizing districts, thereby making the seat-share distribution more variable. When we increase the initial number of voters even more, the trend continues, with the district vote share and total vote share both approaching a point mass at $1/2$.

Going in the other direction, we considered the case where the initial conditions can be rapidly forgotten as districts become swamped by imitated voters (alternatively, the initial number of voters is a small fraction of 1). For example, we recomputed the above

FIGURE 9. District vote (L), total vote (M) and seat (R) distribution when $p = 0.2$: initial conditions 2 : 2.




FIGURE 10. Vote - seat mapping when $p = 0$ (L) and $p = 0.2$ (R), initial conditions 2 : 2.




FIGURE 11. District vote (L), total vote (M) and seat (R) distribution when $p = 0$: initial conditions 1 : 1, $K = 5$





FIGURE 12. District vote (L), total vote (M) and seat (R) distribution when $p = 0.2$: initial conditions 1 : 1, $K = 5$



results when each party has 1 voter in each district initially, but $K = 5$ new voters are added at each step. These are shown in Figures 11–13. When $p = 0$, note the very different vote share distribution but the similarity of the other quantities presented. Also note that when $p = 0.2$ we get results quite similar to the other cases presented above — imitation washes out the effect of initial conditions.

FIGURE 13. Vote - seat mapping when $p = 0$ (L) and $p = 0.2$ (R), initial conditions 1 : 1, $K = 5$




FIGURE 14. District vote (L), total vote (M) and seat (R) distribution when $p = 0$: initial conditions 2 : 1/1 : 2






FIGURE 15. District vote (L), total vote (M) and seat (R) distribution when $p = 0.2$: initial conditions 2 : 1/1 : 2



3.2. Other initial conditions. We next considered initial conditions in which for the northern districts, Party 1 starts with the advantage of two initial voters to Party 2's single voter, while this advantage is reversed in the southern districts. In the absence of inter-district imitation (*i.e.* $p = 0$), this means that the district-level vote distribution is now Beta(2, 1), or triangular (Figure 14, left panel), with each district having probability $\frac{3}{4}$ of being won by its dominant party. The resulting seat-share distribution resembles Figure 4 (left panel), but with less variance.

The polarisation of this electorate along geographical lines can be reduced by increasing the value of p . Imitation across the north-south boundary has the effect of making at least some districts more competitive, and increasing the variance of the seat-share distribution. This is perhaps a step towards a more realistic version of our model: real-world elections typically have many districts in which the mild initial dominance of one party can be overcome by national-level effects.

FIGURE 16. Seat shares for small party under Scenarios (i) – (iii) in Section 3.3 ($p = 0.1$)



We repeated the analysis with initial bias of 3 : 1 instead of 2 : 1, and the results were as expected. For $p = 0$, there is a continuation of the trends in district vote and seat distribution, and a shift in the mean of the total vote distribution but little change in spread. The slope of the seat-vote curve decreases. With $p = 0.2$, the effect of the initial conditions has been largely washed out.

In summary, we have found that increasing p has the effect of concentrating the vote distribution in a single district, leaving the total vote distribution essential unchanged, and spreading out the seat distribution. Furthermore it increases north-south correlation and steepens the seat-vote curve. Changing the initial conditions of the process (but keeping the symmetry between parties overall) has a strong effect on the shape of the vote distribution in a given district, but does not change anything else substantially.

3.3. Modelling third parties. In order to investigate scenarios involving a smaller third party competing with two roughly equal larger ones, we ran our model with the following sets of initial conditions.

- (i) in 80 districts Parties 1,2,3 have respectively 0,2,2 voters; in 10 districts they have 1,2,2 and in 10 districts they have 2,1,1.
- (ii) as above, but in the last 10 districts (the “regional base”) the initial conditions are 3,1,1.
- (iii) in all districts parties 1,2,3 have respectively 1,2,2 voters.

We expect that the small party will fail to become large. When $p = 0$, it cannot extend to any of the first 80 seats in the first two cases. As p increases, the imitation means that although it attracts voters in districts away from its base, it loses more in its base to the other parties. We present graphs on seat distribution for Party 1 under the three scenarios in Figure 16.

4. APPLICATIONS

4.1. The cube law. It has long been observed (see *e.g.* [KS50]) that in two-party electoral systems where each district elects a single representative by simple plurality voting (“First Past the Post”, or FPP), there is a relationship between the fraction x of all the votes and the fraction y of all the districts won by a party along the lines of

$$(1) \quad \frac{y}{1-y} \approx \left(\frac{x}{1-x} \right)^k.$$

When $k = 3$ this is sometimes known as the “cube law”, or “law of the cubic proportion”. For a more comprehensive discussion of this “law” and some alternatives to it, see [Tuf73; Taa73; Taa86; GT79].

TABLE 1. Best fit slope at (50, 50)

| p | initial conditions 1:1 | initial conditions 2:2 | initial conditions 2:1/1:2 | initial conditions 3:1/1:3 |
|-----|------------------------|------------------------|----------------------------|----------------------------|
| 0.0 | 1.48 | 1.67 | 1.52 | 1.25 |
| 0.1 | 2.59 | 2.91 | 2.58 | 2.10 |
| 0.2 | 4.79 | 5.25 | 4.82 | 4.29 |

Taking (1) as an equality gives

$$(2) \quad y = \frac{x^k}{x^k + (1-x)^k} \quad \text{and in particular,} \quad \left. \frac{dy}{dx} \right|_{x=\frac{1}{2}} = k.$$

It is principally this last result that has an empirical basis: for a party winning around 50% of all the votes, each additional 1 percent of vote-share translates into an additional k percent of the parliamentary seats.

The two-party version of our model also exhibits behaviour of this kind, with the value of k depending on the inter-district imitation parameter p . Figure 17 shows our results when initially there is 1 voter of each of 2 parties in each district, and $p = 0.5$. The fitted (red) line corresponds to the formula (2) with $k = 30$. The full curve is visible only in rather unrealistic situations like this one where strong self-reinforcement (*i.e.* large p) makes extreme (“landslide”) election victories for either party possible. The seat-vote mappings we have seen in Figure 6 and Figure 10, along with those derived from real-world elections [GT79, Ch. 1], show only a small part of the full curve, because extreme landslides are very unlikely in practice unless there is substantial correlation between districts.

In general, k increases with p : imitation between districts increases the advantage in seats to a party of gaining extra vote share overall, as can be seen in Figure 6. This seems intuitively reasonable: even “useless” extra voters in districts that their chosen party has no real chance of winning (or no real chance of losing) may still do the party some good if they can influence other voters elsewhere on the electoral map.


Table 1 shows the best fit for k obtained by ordinary linear regression, for various initial conditions and values of p . The value $k = 3$ corresponds then to a value of p slightly more than 0.1. Note that each additional increase of one unit in total vote share gives an increase in seat share that is more than one unit, even when the districts are independent ($p = 0$).

For the simplest independent-district cases, it is possible to do some direct calculations of the model’s seat-vote relationship. Let X_1, \dots, X_N be independent, identically distributed random variables representing a party’s vote shares in each of N identically-sized districts. Assume that the common distribution of the X_i is the uniform distribution; this corresponds to the large- n limit of the initial model depicted in Figure 2. Let $\bar{X}_N = \frac{1}{N} \sum_{i=1}^N X_i$, the total vote share. Then the theoretical expected seat-share y corresponding to a total-vote share x is

$$(3) \quad y = P\left(X_1 > \frac{1}{2} \mid \bar{X}_N = x\right)$$

(since the expected seat share is the same as the probability of winning the first (or any given) district). For $N = 2$, (3) yields the exact result

$$(4) \quad y = \begin{cases} 0 & , \text{if } x \leq \frac{1}{4} \\ 1 - \frac{1}{4x} & , \text{if } \frac{1}{4} \leq x \leq \frac{1}{2} \\ \frac{1}{4(1-x)} & , \text{if } \frac{1}{2} \leq x \leq \frac{3}{4} \\ 1 & , \text{if } x \geq \frac{3}{4}. \end{cases}$$

FIGURE 17. Seats versus votes when $p = 0.5$, initial conditions 1:1

From this, and a similar but lengthier calculation for $N = 3$, we obtain

$$(5) \quad \left. \frac{dy}{dx} \right|_{x=\frac{1}{2}} = \begin{cases} 1 & , \text{if } N = 2 \\ 2 & , \text{if } N = 3. \end{cases}$$


Our simulations suggest that these are extreme cases, with the central slope being close to 1.5 for $N \geq 4$.

4.2. Inter-election swing. In many situations, repeated opinion polling at the level of districts is infeasible. Instead, preference changes at district level are often imputed from the overall national *swing* using a model. The most common are the *uniform swing* model and the *proportional swing* model. These models are used, for example, for election prediction — knowing the exact district-level information from a base election and an estimate of the national swing, we infer the district level information for the next election and hence the election result.

We can test these models against our simulated data as follows. We first run the urn process with 2 parties and 100 districts, starting with 1 voter of each party in each district, until we have 1000000 voters. We then rescale down to 600 voters, keeping the party percentages the same in each district (up to rounding error) and re-run the process. This typically results in a small difference in each district and overall, which gives the local and overall swing. We did this 1000 times and recorded the results from District 1 each time, yielding 1000 data points.

To analyse the data, we plotted the difference between local and national swing against the original vote share in the district. Results are shown in Figure 18. These show clearly that uniform swing fits better for $p = 0$ but proportional swing fits better for $p = 0.1$. Under

FIGURE 18. Local minus national swing versus original vote share in District 1: $p = 0$ (L) and $p = 0.1$ (R)



uniform swing we expect a horizontal line, while under proportional swing a line with positive slope, moving from negative to positive values.

4.3. Retrospective simulations of real elections. We investigated the 2015 UK general election from the point of view of what might have happened according to our model. The approach is as follows. We take the real vote data from the election, for each party in each district. We then downscale this so that each district has 30 voters and the vote shares of parties change little (we achieve this by using the Webster–St-Lagüe allocation method). We then run our urn model with $K = 1$ and $p = 0.08$ until the exact number of voters in each district is achieved (in practice, the process converges quickly enough that it may be terminated once the winner of a district is clear). We performed 100 runs.

Results are summarized in ***.

We then investigated a district-level phenomenon, namely the payoff to a party if it is able to add 1 voter to one downscaled district (effectively an increase of 3 percentage points to their popularity in that district). For each district we ran *** simulations *** random seed ***.


Results ***average number seats, probability of achieving certain number of seats **

5. CONCLUSION

Our results above demonstrate the flexibility and power of the model — with relatively few natural parameters we are able to reproduce a wide range of observed behaviour. In particular with 2 parties, symmetric initial conditions, and $p = 0.1$ we can obtain rather realistic-looking data for FPP elections (compare the results shown in this paper with the seat-vote mappings and other data shown in [GT79], for example, and compare Figure 19 with Figure 1). This allows for a variety of applications to the analysis of electoral systems using plurality ballots. This includes systems based on multi-member districts in addition to the “First Past the Post” setup considered in the present paper. For example, we plan a theoretical test of the tradeoff between decisiveness and proportionality of such systems.

Another use of simulated data generated via the method of this paper is to the analysis of the vote-seat mapping when multi-member districts are involved. Taagepera [Taa86] generalized the basic cube law formula to the multi-member district case, and one could carry out a best fit analysis similar to that in Section 4.1 for that situation.

One possible extension to the model is to remove the symmetry in the way imitation occurs. For example, some districts may be more influential than others. Or, the sphere

FIGURE 19. A simulation with $N = 571$, $m = 2$, $p = 0.1$ 

of influence of a district may extend only locally – to other districts geographically or demographically close to it – rather than to all other districts as assumed in the present paper. In general, an influence structure could be modelled by a graph whose nodes are districts, with imitation probabilities on the (directed) edges. Something like this is probably needed to generate a model of the “vanishing marginals” phenomenon observed most strongly in the United States, in which the districts self-organize (or are exogenously organized) into blocs favouring one party or the other, with very few competitive marginal districts. For models including third parties, another possibility is a regionally-based party with supporters who are less likely than most to imitate voters in other districts, or from other parties.

Another possible extension, which we also leave for later, is to consider more complicated preferences. For example, the model works with an arbitrary number of “colors”, each of which represented a candidate under plurality voting in this article. Instead, each color could be a linear preference order. In that case, however, it is less reasonable to assume that each change from a given color to another is equally likely, and a model that incorporates information about the structure of the preferences may be more useful. For example, a transposition of two adjacent elements of the preference order may be more likely than reversing the order entirely. More generally, we could model sincere preferences directly, as distinct from possibly strategic voting behavior.

Gudgin & Taylor [GT79, Sec. 3.2] develop a probabilistic model for voting in districts that resembles our model in that it considers voters one at a time and has them imitate previously added voters. However in their model, imitation is only of recently added voters and the model does not explicitly use inter-district imitation. In contrast to our model, their parameters depend on the desired population size. We are not aware of further use of this model.

Finally, we note that a variant of our model in which the total voting population of each district remains fixed, and voters periodically change their state (*i.e.* opinion) by imitating others as in the present paper, is essentially a version of the well-known *voter model* ([Lig85]). The standard voter model is defined on the integer lattice and emphasizes very local interactions between voters (each voter is influenced only by a handful of immediate neighbours). The possibility of a more general version defined on a graph, with a voter’s range of immediate influence extending throughout a large “district” and even beyond, has not, to the authors’ knowledge, been explored.

REFERENCES

- [BL94] Sven Berg and Dominique Lepelley. “On probability models in voting theory”. *Statistica Neerlandica* 48.2 (1994), 133–146.
- [BM73] D. Blackwell and J. MacQueen. “Ferguson distributions via Pólya urn schemes”. *Annals of Statistics* 1.2 (1973), 353–355.
- [EP23] F. Eggenberger and G. Pólya. “Über die Statistik vorketteter vorgänge”. *Zeitschrift für Angewandte Mathematik und Mechanik* 3.4 (1923), 279–289.
- [GT79] G. Gudgin and P.J. Taylor. *Seats, votes, and the spatial organisation of elections*. London: Pion, 1979.
- [KS50] M.G. Kendall and A. Stuart. “The law of the cubic proportion in election results”. *British Journal of Sociology* 1 (1950), 183–196.
- [Lig85] T. M. Liggett. *Interacting particle systems*. New York: Springer, 1985.
- [Reg+06] Michel Regenwetter, Bernard Grofman, Anthony Marley, and Ilia Tsetlin. “Behavioral social choice”. *Cambridge University Press* 13 (2006), 58–68.
- [Taa73] Rein Taagepera. “Seats and votes: A generalization of the cube law of elections”. *Social Science Research* 2.3 (1973), 257–275.
- [Taa86] Rein Taagepera. “Reformulating the cube law for proportional representation elections”. *American Political Science Review* 80.02 (1986), 489–504.
- [Tuf73] E. R. Tufte. “The relationship between seats and votes in two-party systems”. *American Political Science Review* 67.2 (1973), 540–554.

UNIVERSITY OF AUCKLAND, AUCKLAND, NEW ZEALAND

E-mail address: g.pritchard@auckland.ac.nz

UNIVERSITY OF AUCKLAND, AUCKLAND, NEW ZEALAND

E-mail address: mcw@cs.auckland.ac.nz